arXiv:1406.2824v1 [cs.SE] 11 Jun 2014

# Some Ideas for Program Verifier Tactics

Gudmund Grov
School of Mathematical and Computer Sciences,
Heriot-Watt University, Edinburgh, UK, G.Grov@hw.ac.uk

**Abstract.** A program verifier is a tool that can be used to verify that a "contract" for a program holds – i.e. given a precondition the program guarantees that a given postcondition holds – by only working at the level of the annotated program. An alternative approach is to use an interactive theorem prover, which enables users to encode common proof patterns as special programs called "tactics". This offers more flexibility than program verifiers, but at the expense of skills required by the user. Here, we add such flexibility to program verifiers by developing "tactics" as a form of program refactoring called DTacs. A formal characterisation and set of examples are given, illustrated with a case study from NASA.

## 1 Introduction

Properties that programs should satisfy are commonly expressed by *contracts* – given a precondition the program guarantees that a given postcondition holds. *Program verifiers* can then be used to verify that a contract is satisfied by automatically generating verification conditions (VCs), which are sent to an underlying theorem prover. Failures to prove VCs will then be highlighted in the text, and the user must then update the code with *auxiliary annotations* to guide the proof. Program verifiers which follow this approach include Spec# [5], VCC [7], Verifast [13], Dafny [14], and the 2014 version of SPARK [3].

An alternative is to translate the annotated programs to an *interactive theorem prover* (ITP) and generate and prove the VCs within this system. However, the disadvantage with this approach is that the user requires expertise in the ITP system as well as the ability to specify and implement the desired program. Many argue that it takes at least six months to become a confident user of an ITP system (see [16]). Such ITP expertise is not required for program verifiers.

An ITP system keeps a stack of open goals, starting with the singleton stack containing the VC. Here, the user can interactively guide the proofs, often in a backwards manner where a 'proof step' involves applying a type of program called a *tactic* to the top of the stack, and pushing any newly generated sub-goals to the stack afterwards [11]. Finding a proof is often a trial-and-error process, which iterates until the stack is empty. A key feature of ITP systems is that common reasoning patterns can be encoded as new tactics, which in most cases are combinations of existing ones. This enables automation of common tasks.

A 'proof step' in a program verifier involves changing, and in most cases adding, auxiliary annotations. This author is not aware of any program verifier with the support for encoding reasoning patterns – and as a result the user will need to manually encode every 'proof step'. As some trial-and-error (i.e. search) may be required, such encoding can be a very tedious and cost-ineffective process. To enable a user to focus time and effort on key decisions, we hypothesise that:

*It is possible to develop "tactics" to automate program verifiers as tactics have been used to automate interactive theorem provers.*

In order to integrate manual proof steps with such "tactic" applications, the user will need work on the annotated program, meaning that a "tactic" application will become a *program transformation*. For example, a typical step is to add an intermediate *assertion* in the code to guide the prover, which means the "tactic" is a transformation which adds that assertion to the code. Other advantages of working at the source code level are that users are not required to have additional expertise, and can inspect the changes a "tactic" has made.

Most modern ITP systems are based on what is known as the LCF approach [11]. Here, the underlying type system ensures that the only trusted part is a small 'kernel' of axioms/rules (together with the type system itself). Thus, if this is sound, then it is guaranteed that (a) the proof is sound; and (b) the original conjecture to be proven is not changed by a tactic. In the "tactics" proposed here a program is transformed to a new program, and the program verifier is re-applied. Thus, (a) is reduced to the correctness of the program verifier, and we are left to focus on (b). Here, one has to ensure that neither the original contract nor the program is changed – only auxiliary annotations can be altered. This can be seen as a special case of a behaviour preserving transformation known as *program refactoring* [10], so we further hypothesise that:

*It is possible to encode program verifier "tactics" as a form of program* refactoring*, to ensure that the contract and program are not changed.*

To address these hypotheses we will focus on the Dafny system [14], described next with an analysis of the proof process. Dafny "tactics", called DTacs, are defined, formalised and exemplified in §3, and applied to a case study from NASA in §4. Finally, we conclude and discuss related and future work in §5.

## 2 Program Verification in Dafny

Dafny [14] is a programming language and program verifier for the .NET platform, developed by Microsoft Research. The language is an imperative object-oriented language, containing both methods and (proper) functions (without side-effects). We may used the term 'methods' for both methods and functions for ease of reading. Fig. 1 illustrates a contract for a Dafny method, where **requires** precedes a precondition and **ensures** precedes a postcondition. By declaring this method as **ghost**, it becomes a specification element, and will not be compiled. In fact, a ghost method can be seen as a "lemma". Variables can also be **ghost** and all variables inside a ghost method are automatically treated as ghost variables. Other annotations include assertions (**assert**) and loop invariants (**invariant**).

Note that programming language elements are often used to support the verification of "lemmas", thus there is close correspondence here with proofs and programs. This is exemplified in Fig. 1. As a result, Dafny has been suggested beyond what is considered to be 'software verification' and also as an alternative to "traditional" theorem provers. One example is found in [15], where Dafny had comparable results to inductive theorem provers.

```
0  ghost method LemmaLength4(n : int)
1    requires n ≥ 0;
2    ensures ∃ xs • length(xs) = n;{
3      if (n = 0){ }else{
4            LemmaLength4(n−1);
5            var xs :| length(xs) = n−1; }}
```

**Fig. 1.** (Almost) a proof of a "lemma" in Dafny (adapted from [16]).

Fig. 1 represents a "lemma" which states that for a given non-negative number n, there exists a list of that length. The proof is by induction. For the step case ($n \neq 0$) the induction hypothesis is applied by making a recursive call to itself with $n-1$ as an argument. This means that we can assume the postcondition of the recursive call, i.e. $\exists\, xs.\ length(xs) = n-1$. Using this property, line 5 creates such a list (:| should be read "such that"). No guidance is given for the base case ($n=0$) as this is trivial to prove for Dafny.

To verify the program, Dafny translates it into an *intermediate verification language* (IVL), called Boogie2 [4]. An IVL can be seen as a layer to ease the process of generating new program verifiers. From Boogie2 a set of VCs are generated and sent to the Z3 SMT solver [18]. If it fails, then the failure is translated back to the Dafny code, via Boogie2. To illustrate, the LemmaLength method above does not actually verify. The problem is that Dafny is not able to determine a witness for the existential in the postcondition. Here, the user needs to give Dafny a hint, achieved by adding the following assertion after line 5:

```
assert length(Cons(1,xs)) = n;
```

Moreover, since LemmaLength is marked to be a **ghost method**, it is only there to aid the verification process. Thus, coming up with this in the first place is likely to have been part of the proof process in [16]. Such "proof steps" in Dafny, will form the basis for a "tactic" in this context. Such steps include:

**Add assertions.** A simple "lemma" can be represented as an assertion in the text, as illustrated above with **assert** length(Cons(1,xs)) = n.

**Create ghost methods.** More involved "lemmas", such as LemmaLength itself that require further proofs, can be captured by a ghost method, where the "lemma" is the postcondition and any required conditions are preconditions.

**Add preconditions.** In certain cases one can restrict the applicability of methods with preconditions. To illustrate, when creating LemmaLength, the $n \geq 0$ precondition could have been added in a second step.

**Proof by cases.** One may need to add code in a ghost method to guide a proof. One common step is to introduce cases, as shown on lines $3-5$ of LemmaLength by using an **if**−**else** statement.

**Add ghost method calls.** One may need to make calls to ghost methods in order to apply "lemmas". This was illustrated for LemmaLength by making a call to itself (to apply the induction hypothesis).

**Add ghost variables.** This is illustrated on line 5 for the LemmaLength method (variables declared in ghost methods are implicitly ghost variables).

**Add postcondition.** One illustration of this is in terms of recursion. A more general lemma is often easier to prove, thus, by strengthening the postcondition the proof may become simpler.

We also consider loop invariant generation an important feature. However, we will not discuss that further here. Our goal is to be able to encode these steps as "tactics" to automate common proof steps for program verifiers.

## 3 DTacs – Dafny Tactics

To illustrate how an LCF tactic works with the stack of sub-goals, consider the stack $[A \wedge B, C]$, and the well-known conjunction introduction (*conj-I*) tactic: $\frac{\vdash A \quad \vdash B}{\vdash A \wedge B}$. Applying *conj-I* will reduce the first sub-goal $A \wedge B$ into the two sub-goals $[A,B]$, which are popped to the stack: $[A,B,C]$.

In a program verifier such as Dafny, failures to prove verification conditions are highlighted in the source code, thus an 'open goal' should here be seen as: the type of failure (e.g. a postcondition), properties of the failure (e.g which postcondition), and the position in the code. The sum of these open goals then correspond to the stack in an LCF prover. A Dafny "tactic", which we call DTac, is a code transformation which creates new (ideally easier to prove) failures when reapplying the program verifier. To illustrate, consider this *MainGoal()* method:

```
method MainGoal() ensures A ∧ B { ... }
```

Assume that the postcondition fails to prove. Here, a DTac which corresponds to the *conj-I* tactic will, when applied, create one **ghost method** for each conjunct and make a call to each of them within the MainGoal() method:

```
ghost method SubGoalA() ensures A { ... }
ghost method SubGoalB() ensures B { ... }
method MainGoal() ensures A ∧ B { SubGoalA(); SubGoalB(); }
```

The program verifier will then solve MainGoal(), but SubGoalA() and SubGoalB() may fail to verify, creating one or two new sub-goals. Thus, the original goal has been reduced into simpler goal(s).

While the LCF approach handles soundness by reducing every proof to the "kernel" of trusted rules [11], we need to restrict the type of transformation allowed, and start with the concept of *program refactoring*. i.e. *"a program transformation which preserves the external behaviour of the original program"* [10]. However, the motivation behind our transformations is to prove that a program satisfies a contract, thus the program should not change at all. Moreover, we cannot allow arbitrary changes to the annotations and contracts. However, as the *conj-I* example showed, a proof is still conducted by changes to the annotations. Thus, we constrain which parts of the code can be changed and which are "final", and define a DTac as follows:

**Definition 1.** *A method is either marked as* 'private' *or* 'public' *– or it has been* 'generated' *by a DTac. We separate between the* 'code' *and* 'specification' *of a program, where the latter is any elements that are not compiled, i.e. contracts, annotations and* **ghost** *code. A DTac is then a refactoring which*

- *cannot change the code for* 'private' *and* 'public' *methods;*
- *cannot strengthen preconditions or weaken postconditions for 'public' methods, but make arbitrary changes to all other specification elements.*

We are interested in proving that the contracts of 'public' methods hold – a 'private' method is simply there as a helper method, and we can change the contract as seen fit. The intuition behind allowing this is that if we e.g. strengthen a precondition, then the caller must prove that it is satisfied when calling it. Anything 'generated' should be seen as part of the proof process and can thus be changed arbitrarily. Note that a user needs to declare which elements are 'private' and which are 'public'. From definition 1 we can see that:

**Theorem 1.** *A DTac does not change the program and preserves the contract of 'public' methods.*

*Proof.* This follows directly from Definition 1.

The following BNF gives the grammar for DTacs:

---

⟨dtac⟩ := ⟨name⟩‘(’ ⟨arg⟩* ‘)’ ‘:=’ ⟨body⟩
⟨body⟩ := ***when*** ⟨prop⟩ ***then*** ⟨trans⟩ | ⟨trans⟩
⟨trans⟩ := ⟨code⟩ $\Longrightarrow$ ⟨code⟩ [ ⟨inst⟩ ] | ***match*** ⟨code⟩ [ ⟨inst⟩ ]
| ⟨trans⟩ ‘;’ ⟨trans⟩ | ⟨name⟩‘(’ ⟨arg⟩* ‘)’ [ ⟨inst⟩ ]
| ***or***(⟨trans⟩,⟨trans⟩)
⟨inst⟩ := ‘[’(⟨pos⟩ | ‘?’⟨name⟩ ‘:=’⟨code⟩)*‘]‘
⟨pos⟩ := ‘@’⟨name⟩ | ***line***‘(’⟨nat⟩*‘)’ | ***up***‘(’⟨pos⟩‘)’ | ***down***‘(’⟨pos⟩‘)’

---

‘$*$’ represents repetition, ‘$[-]$’ for optional, and ‘|’ for alternatives. ⟨code⟩ is the code and annotation of Dafny, populated with variables (preceded by ‘?’) and a special ***rewrite*** function. This function is overloaded to work for: a single rewrite rule ***rewrite***$(l \to r, c)$; a triple ***rewrite***$(l, r, c)$ or a list ***rewrite***$([l_1, \cdots, l_n], [r_1, \cdots, r_n], c)$, which is applied pairwise ***rewrite***$(l_i, r_i, c)$. In all cases the code $c$ is rewritten. We refer to standard rewriting literature for details [2]. For readability, we will also use ellipses ‘...’ for parts of the code that can be ignored instead of variables. $\langle code_1 \rangle \Longrightarrow \langle code_2 \rangle$ represents a rewrite rule, and thus $\langle code_2 \rangle$ cannot contain any variables that are not present in $\langle code_1 \rangle$. ‘;’ is sequential composition, while ***when*** and ***then*** are used to express conditions. ***match*** is used to check for particular code patterns and instantiate variables, typically used with a sequentially composed DTacs. ***or*** is used to express choice.

We assume the presence of an *environment* E, which contains binding of variables and positions which are preceeded by ‘@’. ⟨prop⟩ is a predicate on the underlying environment and the source code. In order to have a simpler model of the composition of DTacs, the variables bindings from the first are removed before applying the second DTac. This introduces a form of “locality” of the namespaces: e.g. for $dt_1; dt_2$, if both binds say variable $?x$ then these can be seen as separate variables. To carry such bindings between DTacs an instantiation environment $\langle inst \rangle$ can be provided. Here, variables can be bound by $[?x := e]$, where $e$ uses the existing environment. To illustrate in $dt_1; dt_2[?x :=?x]$, $?x$ in $dt_2$ is defined to be the same $?x$ as in $dt_1$. A second usage of the instantiation environment is to constrain the position in the source code where a DTac can be applied – either by giving reference to special comment preceded by ‘@’, or line numbers. If $dt[@p1, @p2]$, then $dt$ can only be applied at one of these positions. Inspired by Huet's *Zippers* [12], we can also move ***up*** and ***down*** a statement

relative to a given position. Note that ***up*/*down*** are only applicable 'within' a method – they will fail if attempting to e.g. move down when at the last statement of the method. Positions are very useful when composing DTacs of the method. Recursive application of DTacs is prohibited. Fig. 3 contains several example DTacs. Given source code $c$ and a DTac $dtac$, $\Downarrow$ refactors $c$ into $c'$:

$$\frac{E \in \underline{\boldsymbol{pv}}(c) \qquad \langle E, c, dtac\rangle \Downarrow_{dt} \langle _, c'\rangle \qquad \vdash c'}{\langle c, dtac\rangle \Downarrow c'}$$

$\underline{\boldsymbol{pv}}$(c) can be seen as the verification condition generator, and returns a set of environments, each containing the bindings: ?$error$ for the type of error; ?$err_arg$ for the property that failed; and $@err_pos$ for the position of the error. $\Downarrow_{dt}$ then evaluates the DTac and code under this environment. The new source has to be type correct, and we assume the presence of a typing relation $\vdash$.

SEQ
$$\frac{\langle E, c, dt_1\rangle \Downarrow_{dt} \langle E'', c''\rangle \quad \langle E'', c'', dt_2\rangle \Downarrow_{dt} \langle E', c'\rangle}{\langle E, c, dt_1; dt_2\rangle \Downarrow_{dt} \langle E', c'\rangle}$$

DTAC
$$\frac{dtac \notin \{\mathbf{or}(_,_), (_\ ;\ _)\} \quad \langle E, c, dtac\rangle \Downarrow_{m} E'' \quad \langle E'', c, dtac\rangle \Downarrow_{app} \langle E', c'\rangle}{\langle E, c, dtac\rangle \Downarrow_{dt} \langle E', c'\rangle}$$

OR1
$$\frac{\langle E, c, dt_1\rangle \Downarrow_{dt} \langle E', c'\rangle}{\langle E, c, \boldsymbol{or}(dt_1, dt_2)\rangle \Downarrow_{dt} \langle E', c'\rangle}$$

OR2
$$\frac{\langle E, c, dt_2\rangle \Downarrow_{dt} \langle E', c'\rangle}{\langle E, c, \boldsymbol{or}(dt_1, dt_2)\rangle \Downarrow_{dt} \langle E', c'\rangle}$$

MATCH1
$$\frac{\langle E, c, c\rangle \Downarrow_{m} E' \quad \langle E', c\rangle \models p}{\langle E, c, \boldsymbol{when}\ p\ \boldsymbol{then}\ c\rangle \Downarrow_{m} E'}$$

MATCH2
$$\frac{E' \in \underline{\boldsymbol{pmatch}}(\underline{\boldsymbol{flush}}(E) \cup \underline{\boldsymbol{inst}}(inst, E), c, c_1)}{\langle E, c, c_1 \Longrightarrow c_2\ inst\rangle \Downarrow_{m} E'}$$

MATCH3
$$\frac{E' \in \underline{\boldsymbol{pmatch}}(\underline{\boldsymbol{flush}}(E) \cup \underline{\boldsymbol{inst}}(inst, E), c, m)}{\langle E, c, \boldsymbol{match}\ m\ inst\rangle \Downarrow_{m} E'}$$

MATCH4
$$\frac{\langle E, c, \underline{\boldsymbol{unfold}}(dt(args))\ inst\rangle \Downarrow_{m} E'}{\langle E, c, dt(args)\ inst\rangle \Downarrow_{m} E'}$$

APP1
$$\frac{c' \in \underline{\boldsymbol{apply}}(E, c_1 \to c_2, c) \quad E' = E[@pos, @start, @end := @m, @s, @e]}{\langle E, c, c_1 \Longrightarrow c_2\ inst\rangle \Downarrow_{app} \langle E', c\rangle}$$

APP2
$$\frac{E' = E[@pos, @start, @end := @m, @s, @e]}{\langle E, c, \boldsymbol{match}\ m\ inst\rangle \Downarrow_{app} \langle E', c\rangle}$$

APP3
$$\frac{\langle E, c, \underline{\boldsymbol{unfold}}(dt(args))\ inst\rangle \Downarrow_{app} \langle E', c'\rangle}{\langle E, c, dt(args)\ inst\rangle \Downarrow_{app} \langle E', c'\rangle}$$

APP3
$$\frac{\langle E, c, c\rangle \Downarrow_{app} \langle E', c'\rangle}{\langle E, c, \boldsymbol{when}\ p\ \boldsymbol{then}\ c\rangle \Downarrow_{app} \langle E', c'\rangle}$$

**Fig. 2.** Evaluation semantics for DTacs.

Fig. 2 defines the $\Downarrow_{dt}$ relation. First sequences, and ***or*** have the obvious semantics. Other DTacs have two phases: a *matching* phase followed by an *application* phase.

During the *matching phase*, if the DTac is conditional (***when*-*then***-), the body is matched then the precondition is checked (by $\models$ whose definition is omitted). $\underline{\boldsymbol{pmatch}}$ takes an environment, code and a code pattern, and returns a set where each element is an environment containing updated with all the matches. In addition to the variables the bindings from the pattern, '?pre' is bound to the preconditions, '?post' to the postconditions, '?meth' to the name

of the method the code is within, and '?arg' to the argument of the method. If there are e.g. many preconditions, then it is used non-deterministically to refer to any of them. Moreover, '@s','@e' and'@m' become bound to the full, start and end positions. ***inst***$(inst, E)$ instantiates the variables in $inst$ using $E$, while ***flush*** removes all binding created by a pattern (e.g. '?pre', '@pos', and '?error' are not removed). ***unfold*** unfolds the given DTac and for simplicity we assume that it has access to all of the definitions.

In the *application phase* ***apply***$(E, r, c)$ represents the application of a rule $r$ under the environment $E$ to the code $c$. Here, '@m' is used to apply it to the same position as the match from the matching phase. For ***pmatch*** and ***apply*** we refer to standard rewriting literature, such as [2], and omit the detail.

From $\Downarrow$ we can see that the composition of DTacs preserves the DTac validity;

**Theorem 2.** *If $dt_1$ and $dt_2$ are DTacs then $dt_1$ ; $dt_2$ and* ***or***$(dt_1, dt_2)$ *are DTacs.*

*Proof.* ***or***$(dt_1, dt_2)$ follows directly from $\Downarrow_{dt}$ and the assumption as either $dt_1$ or $dt_2$ is applied. Following $\Downarrow_{dt}$, $dt_1$ ; $dt_2$ first applies $dt_1$, giving $c''$, which by the assumption has only made valid changes following Definition 1. $dt_2$ is then applied to $c''$, which by the assumption has only made changes following Definition 1. Thus, the $dt_1$ ; $dt_2$ is a DTac following Definition 1.

The first six DTacs of Fig. 3 are adaptations of standard introduction and elimination tactics. Note that $\Downarrow$ ensures that the resulting code is well-typed by $\vdash$. *post-to-assert* turns a postcondition into an assertion at the end of the method position (achieved by @end), and is normally applied due to a failed postcondition. The *assert-to-pre* DTac turns an assertion in the beginning of a method into a precondition, and can be seen as a way of moving the work of proving $?P$ to the caller. Similarly, *assert-to-post* moves an assertion to a postcondition of the method preceding it. Here, ***rewrite*** is used to adapt the arguments to the new context. E.g. if the assertion was `x+1 < y` with method call `?m(x+1,y)` for **method** `?m(a,b)`, then the postcondition would become `a < b`. *assert-rewr* applies a given rewrite rule to an assertion. *assert-up* moves an assertion upwards in the code passing a statement. It is moved into both branches of an **if**$-$**else** statement, and updated with the expression in case of an assignment (both achieved by ***or***). A failed postcondition of a method could be due to a failed postcondition of a nested method call. In this case, the postcondition needs to be "copied" to the method called, achieved by *post-to-post*. *pre-to-assert* and *null-to-assert* are examples of DTacs triggered by a particular type of failure, and in both cases an assertion is added just before the failure. Given a fresh variable and a predicate, *pred-var-I* is used to introduce a new ghost variable, such that the predicate holds. A special use of this is to eliminate an existential quantifier, as shown in the *ex-E* DTac. *case-I* introduces a case split, using an **if**$-$**else** statement, whilst *call-I* inserts a call to a ghost method. *IH-I* is a special case of *call-I*, where a recursive call is made to itself, with the argument decremented by one, at the @case2 position. In addition to those in Fig. 3, five more DTacs are required for the case study. The definitions of these have been omitted for space reasons. *assert-down* moves an assertion down. *assert-conj-I* splits a conjunction

*assert-I*(P) := ☐ ⟹ `assert ?P; /*@ass*/`

*post-I*(P) := `method ?m(..) ... { ... }` ⟹ `method ?m(..) ... ensures ?P { ... }`

*pre-I*(P) := ***when*** ?*m* is not public ***then*** `method ?m(..) ...` ⟹ `method ?m(..) requires ?P...`

*assert-E*() := `assert P;` ⟹ `/* @ass */`

*pre-E*() := `method ?m(..) requires ?P...` ⟹ `method ?m(..) ...`

*post-E*() := ***when*** ?*m* is not public ***then*** `method ?m(..) ... ensures ?P { ... }` ⟹ `method ?m(..) ... { ... }`

*post-to-assert*() := ***match*** `method ?m(..) ... ensures ?P...` ; *assert-I*(?P)[@end,?meth := ?m]

*assert-to-pre*() := *assert-E*()[@start]; *pre-I*(?P)[?meth := ?meth]

*assert-to-post*() := ***match*** `?m(?xs); assert ?P;` ; *assert-E*()[@pos,?P := ?P,?m := ?m] ;
***match*** `method ?m(?ys)` [?P := ?P,?m := ?m,?xs := ?xs];
*post-I*(***rewrite***(?*xs*, ?*ys*, ?*P*))[?m := ?m]

*assert-rewr*(R) := *assert-E*();*assert-I*(***rewrite***(R,?P)[@pos]

*assert-up1*() := *assert-E*();*assert-I*(?P)[***up***(@pos)]

*assert-up2*() := ***match*** `?x := ?e; assert ?P;` ; *assert-E*()[@pos,?P:=?P, ?x:=?x,?e:=?e];*assert-I*(***rewrite***(?x,?e,?P)[***up***(@pos)]

*assert-up3*() := `if(..){..}else{..} assert ?P;` ⟹ `if(..){../*@ass1/*assert ?P;} else {../*@ass1/*assert ?P;}`

*assert-up*() := ***or***(*assert-up3*(),***or***(*assert-up2()*,*assert-up1*()))

*post-to-post*() := ***match*** `method ?m1(..) ... ensures ?P; { ... ?m2(?xs); ... }` ;
***match*** `method ?m2(?ys)` [?m1:=?m1,?m2 := ?m2, ?xs := ?xs, ?P := ?P] ; *post-I*(***rewrite***(?*xs*, ?*ys*, ?*P*))[?meth := ?m2]

*pre-to-assert*() := ***when*** ?error = "A precondition for this call might not hold" ***then*** *assert-I*(?*err_arg*)[***up***(@err_pos)]

*null-to-assert*() := ***when*** ?error = "target object may be null" ***then*** *assert-I*(?*err_arg* ≠ **null**)[***up***(@err_pos)]

*pred-var-I*(v,P) := ☐ ⟹ `ghost var v :| P /*@gv*/`

*ex-E*(x,P) := ***when*** $P = \exists\ ?y.?P'$ ***then*** *pred-var-I*(x,***rewrite***(?y,x,?$P'$))

*case-I*(cond) := ***when*** ?meth is generated ***then*** ☐ ⟹ `if (cond){ /*@case1*/}else{/*@case2*/ }`

*call-I*(m,args) := ***when*** *m* is ghost ***then*** ☐ ⟹ `m(args); /* @call */`

*IH-I*() := *call-I*(?meth,?arg - 1)[@case2]

***Correctness:*** *assert-I* and *assert-up3* only adds a specification element. *post-I* only strengthens the postcondition, while *pre-E* only weakens a precondition. *pre-I* and *post-E* have preconditions ruling out application to public methods. *assert-E* only removes an assertion *pred-var-I* only introduces ghost code, while *case-I* and *call-I* are restricted to ghost code. The remaining are just composition or special cases of other DTacs and are thus correct.

**Fig. 3.** DTac examples

in an assertion into one assertion for each conjunct. *assert-up-ctxt()* behaves as *assert-up*, but preserves "context" when within an **if**(−**else**) block. E.g. assertion P and condition C, becomes C ⟹ P when moving out of the **if** block and ¬C ⟹ P when moving out of the **else** branch. *assert-strengthen()* removes a condition (conjunct) from a conditional assertion. Finally, *assert-comb1()* is a very low-level DTac which turns **assert** P ⟹ Q; **assert** P ⟹ R into **assert** Q ⟹ R.

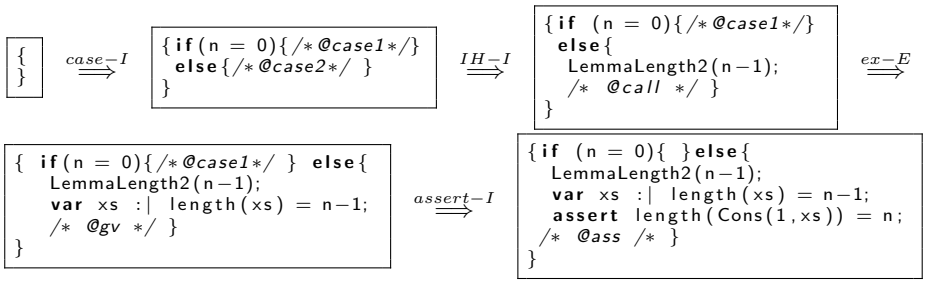


**Fig. 4.** Proof of the LemmaLength method using DTacs.

Fig. 4 is a proof of the LemmaLength **ghost method**. The starting point is the method of Fig. 1, but with an empty body. This fails to prove and the first step applies *case-I*()[?meth := LemmaLength]. *IH-I* is then used to apply the "induction hypothesis" where the existential is instantiated by *ex-E*(xs,∃ xs. length(xs) =n)[@call]. The proof is completed by giving the witness for the postcondition as an assertion: *assert-I*(length(Cons(1,xs)) =n)[@gv]. The next section contains a larger case study and more realistic applications of DTacs.

## 4 A SAFER System with DTacs

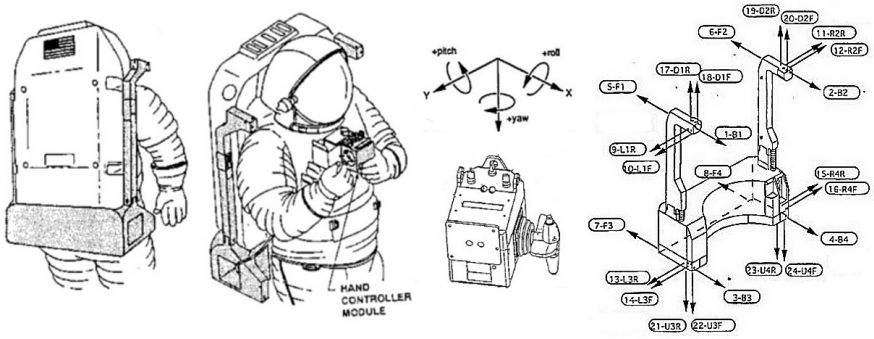


**Fig. 5.** Left-to-right: SAFER deployed on an astronaut; the six degrees of movement axes (top); the hand controller (bottom); the mounting of the thrusters. (source: [19])

Simplified Aid For Eva Rescue (SAFER) is a lightweight backpack propulsion system developed by NASA, and intended to provide self-rescue capabilities for astronauts outside the spacecraft in space. It enabes movement along and around all three axes. At frequent intervals, the software system reads commands from a *hand-controller module* (HCM), and various sensors, and selects which of the

mounted 24 thrusters should be fired (by opening a vent releasing $GN_2$ gas). In addition to the commands, the system contains an *automatic attitude hold* (AAH), which attempts to nullify rotation. The system is shown in Fig. 5.

Several properties of a smaller version of the system have been verified using PVS [19] and validated by testing using VDM-SL [1]. Here, a subset of this system is implemented in Dafny, and DTacs are used to ensure freedom of **null** references, together with a key functional property: only four thrusters can be simultaneously fired[1]. To encode this the following types are used:

```
datatype cmd      = NEG | ZERO | POS;
datatype thruster = B1 | B2 | B3 | B4 | F1 | ...;
datatype switch   = TRAN | ROT;
class T_CMD  { var X : cmd; var Y : cmd; var Z : cmd; ... }
class R_CMD  { var roll : cmd; var pitch : cmd; var yaw : cmd; ..}
class R_PRED { var roll : bool; var pitch : bool; var yaw : bool; ..}
class SD_CMD { var tran : T_CMD; var rot : R_CMD; ... }
```

cmd describes the orientation w.r.t the six axes, thruster is the name of each thruster, while switch captures the mode of the HCM (see Fig. 5). The class T_CMD holds the "transitional" movements, R_CMD "rotational" movement, and SD_CMD combines them. R_PRED is a predicate for each rotation axis.

The HCM has a mode switch, and a stick that can be moved in four directions (along the three axes and twisting). In each step, these values are read from the HCM and given as input to the control method (note that $|s|$ is the length of sequence s):

```
0 method control(vert:cmd, horiz:cmd, trans:cmd, twist:cmd, mode:switch)
1  returns (thrusters : seq<thruster>) modifies this; {
2  ensures |thrusters| ≤ 4;
3     var aah_cmd := AAH_cmd();
4     var grip_cmd := grip_command(vert, horiz, trans, twist, mode);
5     var cmds := integrated_cmds (grip_cmd, aah_cmd);
6     thrusters := selected_thrusters(cmds); }
```

This method returns the set of thrusters to fire, and the key property is a poscondition of the method. The additional sensor readings are only required by the AAH. The four-thruster property is independent of the AAH, and has therefore been omitted from the program. Thus, we ignore the other sensors. However, we still need some inputs from the AAH, which are left undefined:

```
method AAH_cmd () returns (res : R_CMD) ensures res ≠ null;
method AAH_ignore_HCM () returns (res : R_PRED) ensures res ≠ null;
predicate method AAH_all_axis_off()
```

AAH_cmd is the rotation command from the AAH. AAH_ignore_HCM is used in certain cases for the AAH to override rotational commands from the HCM. Finally, AAH_all_axis_off is true if the AAH is off for all axes. Since these are left unimplemented, the property we prove is true regardless of what is returned from these methods. These three methods, together with control, are the only 'public' methods. The remaining, discussed next, are 'private', meaning DTacs are free to alter their contracts.

[1] Various snapshots from verifying these properties in Dafny 1.7.0 can be found at `https://sites.google.com/site/gudmundgrov/research/DafnySAFER.zip`.

On line 4 of control, grip_command(vert,horiz, trans, twist, mode) turns the inputs into a SD_CMD object. If the mode is TRAN then these are mainly transitional commands, otherwise they are mainly rotational commands. Regardless of the mode, an X and pitch command will always be issued. On line 5 of control, the commands from the HCM and AAH are combined into a single command by integrated_cmds. Here, if all AAH axes are off, then the command becomes the command from the HCM. Here, rotational commands have priority: if such are present, then only rotational commands are issued. If not, a transition command on one axis is issued, with the priority: $X > Y > Z$. If it is not the case that all AAH axes are off, and both AAH and HCM rotational commands are present, then HCM has priority unless AAH_ignore_HCM specifies that the HCM should be ignored for that particular axis.

The key logic for the four thruster property is within the selected_thrusters method, applied at the end of control. The encoding here follows directly from [19]. Firstly, the thrusters are divided into those mounted in a Back or Front (BF) position, and those in a Left-Right-Up-Down (LRUD) position. Moreover, "mandatory" and "optional" thrusters are separated, creating four Dafny functions – taking three axis commands and returning a sequence of thrusters. E.g.

```
function method BF_optional(A:cmd,B:cmd,C:cmd) : seq<thruster> {
  match A case NEG ⇒ (match B case NEG ⇒ (match C case NEG ⇒ [B2,B3]...
```

returns the optional BF thrusters. Next, using these functions, methods for the BF and LRUD as a whole are created, exemplified for BF:

```
method BF (A:cmd,B:cmd,C:cmd) returns (man:seq<thruster>,opt:seq<thruster>){
 man := BF_mandatory(A,B,C); opt := BF_optional(A,B,C); }
```

The selected_thrusters method then combines these depending on whether or not X, pitch and yaw commands are present:

```
method selected_thrusters (comb : SD_CMD) returns (thrusters : seq<thruster>){
 var bf_main, bf_opt := BF(comb.tran.X,comb.rot.pitch,comb.rot.yaw);
 var lrud_main, lrud_opt := LRUD(comb.tran.Y,comb.tran.Z,comb.rot.roll);
 if (comb.tran.X = ZERO){
   if (comb.rot.pitch = ZERO ∧ comb.rot.yaw = ZERO){
      thrusters := bf_opt + bf_main + lrud_opt + lrud_main;
   } else{  thrusters := bf_opt + bf_main + lrud_main; }
 } else {
    if (comb.rot.pitch = ZERO ∧ comb.rot.yaw = ZERO){
     thrusters := bf_main + lrud_opt + lrud_main;
    }else{ thrusters := bf_main + lrud_main; }}}
```

**Handling null references** Dafny gives 16 errors when trying to verify the above example, and 15 of them are "target object may be null" thus we need to rule out the possibility of **null** references. We will illustrate a common strategy to prove such goals with one such failure in the selected_thrusters method. Here, the problem is the use of comb in the call to BF. First, the *null-to-assert* DTac is applied, resulting in the following update to the code:

```
 assert comb ≠ null;
 var bf_main, bf_opt := BF(comb.tran.X,comb.rot.pitch,comb.rot.yaw);
```

This assertion fails and the *assert-to-pre* DTac is applied, which moves this assertion to a precondition:

```
method selected_thrusters ... requires comb ≠ null; ...
```

Now this condition does not hold in the call to selected_thrusters by the control method, raising the "a precondition for this call might not hold" error. By applying the *pre-to-assert* DTac, control is updated as follows:

```
var cmds := integrated_cmds (grip_cmd, aah_cmd, active_axis, ignore_HCM);
assert cmds ≠ null; thrusters := selected_thrusters(cmds);
```

This assertion fails to verify. We then apply *assert-to-post* which moves this assertion to the postcondition of integrated_cmds:

```
method integrated_cmds ... returns (comb : SD_CMD) ensures comb ≠ null;
```

This completes the proof, and all such errors follow a similar approach. The only difference is that in certain cases, the final postcondition cannot be verified. In those cases the *post-to-assert* DTac, which creates an assertion with the postcondition at the end of the branch, is applied. This assertion is then moved upwards by *assert-up*, and eventually becomes a precondition that is not met by the caller. This is resolved by the *pre-to-assert* DTac. At this point the above approach is followed. In total *null-to-assert* is applied 10 times; *assert-to-pre* 11 times; *pre-to-assert* 5 times; *assert-to-post* 10 times; *assert-up* 19 times; and *post-to-assert* 4 times.

**Verifying the four thruster property** After handling all the "**null** errors", the only open goal is the main property: the $|\,\mathsf{thrusters}\,| \leq 4$ postcondition of control. The main logic for this computation is within the selected_thrusters method so the first step is to move the goal there: *post-to-assert* moves it to an assertion in control, and then *assert-to-post* moves it to a postcondition of selected_thrusters. *post-to-assert* then moves the goal into the code, while *assert-up* (several) times, reduces this to 4 goals, one for each branch.

First, the last branch is addressed. Here, *assert-up* moves the assertion "over" the assignment, replacing thrusters with the assignment expression, resulting in:

```
assert |bf_main + lrud_main| ≤ 4; thrusters := bf_main + lrud_main;
```

*assert-up* is then applied until the assertion is just below the call to BF. The following rewrite rules will be required in the remainder:

$$|\,?x + ?y\,| \leq ?n \longrightarrow |\,?x\,| \leq (?n\ /\ 2) \ \wedge\ |\,?y\,| \leq (?n\ /\ 2) \tag{1}$$

$$|\,?x + ?y\,| \leq ?n \longrightarrow |\,?x\,| \leq ?n \ \wedge\ |\,?y\,| = 0 \tag{2}$$

$$|\,?x + ?y\,| \leq ?n \longrightarrow |\,?x\,| = 0 \ \wedge\ |\,?y\,| \leq ?n \tag{3}$$

We first apply *assert-rewr*((1)) followed by *assert-conj-I*, resulting in:

```
assert |lrud_main| ≤ 4/2; assert |bf_main| ≤ 4/2;
```

This is followed by *assert-to-post*, and then *move-up* followed by *assert-to-post*, which moves these into a postcondition $|\mathsf{man}| \leq 4/2$ for both LRUD and BF, completing the first goal.

Next, we address the property at the first branch, by applying *assert-up*:

```
assert |bf_opt + bf_main + lrud_opt + lrud_main| ≤ 4;
thrusters := bf_opt + bf_main + lrud_opt + lrud_main;
```

Then *assert-up-ctxt* twice, creates this assertion just after the call to LRUD:

```
assert comb.tran.X = ZERO ∧ comb.rot.pitch = ZERO ∧ comb.rot.yaw = ZERO
                    ⟹ |bf_opt + bf_main + lrud_opt + lrud_main| ≤ 4;
```

A combination of *assert-rewr*((1)), *assert-rewr*((2)), *assert-rewr*((3)) and *assert-conj-I* results in the following four assertions (with the same hypothesis as above): |bf_opt| =0; |bf_main| =0; |lrud_opt| ≤4/2; and |lrud_main| ≤4/2. |lrud_main| ≤4/2 is already handled by the postcondition of LRUD introduced by the first goal. Of the remaining, the first two goals are verified by *move-up* followed by *assert-to-post* for BF, while the third one is verified by *assert-to-post* for LRUD, after *assert-strengthen* has been applied trice to remove the conditions.

The same approach is applied to the third goal, where the condition is comb.tran.X =ZERO ∧ ¬(comb.rot.pitch =ZERO ∧ comb.rot.yaw =ZERO), and conditional assertions have the conclusions |bf_opt| =0, |bf_main| ≤2 and |lrud_main| ≤2. Only the first of these is not provable, which becomes a postcondition for BF by *move-up* followed by *assert-to-post*.

We follow the same approach for the final goal, which has the condition comb.tran.X ≠ ZERO ∧ comb.rot.pitch =ZERO ∧ comb.rot.yaw =ZERO, and the three goals |lrud_opt| =0, |bf_main| ≤2 and |lrud_main| ≤2. In this case, only the first goal is not proven. This property relies on the fact that there will only be one transition command at a time, thus, since $X$ is not ZERO, that means that $Y$ and $Z$ are. We thus add the following precondition to selected_thrusters

```
requires comb.tran.X ≠ ZERO ⟹ comb.tran.Y = ZERO ∧ comb.tran.Z = ZERO;
```

by the *pre-I* DTac. As a result of this, the precondition to the selected_thrusters call in control is violated. We then apply the *pre-to-assert* DTac, which adds the assertion before the call, followed by the *assert-to-post* that adds this as a postcondition to integrated_commands. This is then proven by a combination of *post-to-assert* and *assert-to-post* DTacs to auxiliary methods (not given here).

The newly added precondition is moved into an assertion by *pre-to-assert*. Then, several *assert-down* DTacs are applied so that it is moved just before the failing |lrud_opt| =0 assertion:

```
assert comb.tran.X ≠ ZERO ⟹ comb.tran.Y = ZERO ∧ comb.tran.Z = ZERO;
assert comb.tran.X ≠ ZERO... ⟹ |lrud_opt| = 0
```

The *assert-comb1* DTac combines these into

```
assert comb.tran.Y = ZERO ∧ comb.tran.Z = ZERO... ⟹ |lrud_opt| = 0
```

*assert-strengthen* (twice) removes pitch and yaw from the precondition (as these are not used by LRUD):

```
assert comb.tran.Y = ZERO ∧ comb.tran.Z = ZERO ⟹ |lrud_opt| = 0;
```

We then apply the *assert-to-post* for LRUD, which completes the proof. In total, 57 DTac applications were required for the "four thruster property".

## 5 Conclusion, Related & Future Work

By developing DTacs and showing a detailed account of their usage to a relevant case study the hypothesis – '*it is possible to automate proofs by a notion of "tactics" for program verifiers as a form of refactoring*' – has been validated. However, understanding the degree of automation will require further examples, particularly those with iteration, and importantly, implementation. As Dafny is open source, the implementation of DTacs can exploit this existing code base, and represent DTacs as transformations on the abstract syntax tree. Such evaluation is likely to generalise, and hopefully improve, the DTacs derived here, giving some indication of how to prune the search. The implementation of SAFER in Dafny was based upon previous embeddings in PVS [19] and VDM-SL [1]. SAFER has here been used as a proof of concept for the feasibility of DTacs, whilst [19,1] uses existing tools and approaches. A direct comparison between the approaches would therefore be unfair at this point.

DTacs are inspired by *tactics* for ITP systems, originating from the LCF system [11]. Most tactic languages are still based on this paradigm. To a less formal extent, refactoring also follows a "trusted kernel" idea, where each refactoring should just make a small change to the code as this is easier to analyse [10]. Larger refactorings then become compositions of smaller ones. DTacs adopt the same approach, as most DTacs only make small changes, and the majority are special cases and/or compositions of smaller DTacs. Composition and specialisation are supported by sequential and branching combinators, and the use of arguments and partial instantiations. Other common combinators, such as "orelse", "try" and "repeat" are likely to be added in the future.

The proofs conducted in the case-study and the LemmaLength example are rather low-level and tedious at times. This is a consequence of working at the lowest "kernel" level. Future work would include the development of higher level strategies which combine low-level rules. Two examples illustrating that this should be possible have been seen: (1) the "null reference" errors all followed the same strategy; and (2) three out of four of the "four thruster" goals followed the same overall strategy. Moreover, the last goal only deviated by requiring an additional precondition.

Chen [6] describes a simple Dijkstra's guarded-command style imperative programming language, including both "normal" execution statements as well as special verification statements. The verification statements can then be used to implement e.g. type checkers, abstract interpreters or contracts for the execution statements. This approach is much more general than DTacs, and whilst verification statements may be used to implement a similar notion of tactics, this is not discussed in [6]. Moreover, it is not clear how such approach will work with an existing programming language and program verifier.

In [20], refactoring of *proof scripts* (for ITP systems) is addressed. Whilst it has similarities to DTacs in the sense that both target proofs, it deviates as it is used to improve existing proofs – albeit, it could be applied to generalise proofs into more general tactics. Further, DTacs deviate from the case-based reasoning used in [17] where specifications are carried between programs. DTacs provides a

mechanism for encoding generic strategies, which could be used to prove both the source and target programs in such a setting. A different approach to refactoring, is to "optimise" the compilation into the IVL, as was done with the 'induction tactic' for Dafny [15]. It is important to note the difference between DTacs and the vast number of static analysers such as abstract interpretation [8]. There, the goal is to discover properties of programs, which could be utilised to prove properties. DTacs provide a way to enable users to encode their strategies when the proof is beyond such methods. However, in the future it would be interesting to see how such work could be utilised in a DTac. Finally, cases of the *assert-up* DTac bear resemblance to a *predicate transformer* [9] – it should therefore be possible to encode other predicate transformers as DTacs in the same way.

This paper demonstrates that the use of "tactics" for program verifiers is feasible and provides a firm foundation for further research in this area.

## References


1. S. Agerholm and P. G. Larsen. Modeling and validating SAFER in VDM-SL. In *4th NASA FM Workshop*. NASA, September 1997.
2. F. Baader and T. Nipkow. *Term Rewriting and all that*. Cambridge University Press, 1998.
3. J. Barnes. *High Integrity Ada: The SPARK Approach*. Addison-Wesley, 1997.
4. M. Barnett, B-Y E. Chang, Robert DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *FMCO*, volume 4111 of *LNCS*, pages 364–387. Springer, 2006.
5. M Barnett, K. R. M. Leino, and W. Schulte. The Spec# programming system: an overview. In *CASSIS*, volume 3362 of *LNCS*, pages 49–69. Springer, 2005.
6. Yifeng Chen. Programmable Verifiers in Imperative Programming. In *Unifying Theories of Programming*, pages 172–187. Springer, 2010.
7. E. Cohen, M. Dahlweid, M. Hillebrand, D. Leinenbach, M. Moskal, T. Santen, W. Schulte, and S. Tobies. VCC: A practical system for verifying concurrent C. In *TPHOL*, volume 5674 of *LNCS*, pages 23–42. Springer, 2009.
8. Patrick Cousot and Radhia Cousot. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *POPL*, pages 238–252, January 1977.
9. E. W. Dijkstra. Guarded commands, nondeterminacy and formal derivation of programs. *CACM*, 18(8):453–457, August 1975.
10. M. Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley, August 1999.
11. M. J. Gordon, R. Milner, and C. P. Wadsworth. *Edinburgh LCF*. Springer, 1979.
12. G. Huet. The Zipper. *J. of Functional Programming*, 7(5):549–554, 1997.
13. B. Jacobs, J. Smans, and F. Piessens. A quick tour of the verifast program verifier. In *APLAS*, volume 6461 of *LNCS*, pages 304–311. Springer, 2010.
14. K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *LPAR*, volume 6355 of *LNCS*, pages 348–370. Springer-Verlag, 2010.
15. K. R. M. Leino. Automating induction with an SMT solver. In *VMCAI*, volume 7148 of *LNCS*, pages 315–331. Springer, 2012.
16. K. R. M. Leino and N. Polikarpova. Verified Calculations. In *VSTTE*, 2013.
17. Pitu M., Grijincu D., Li P., Saleem A., R Monahan, and D.P. O'Donoghue. Aris : Analogical Reasoning for reuse of Implementation & Specification. In *AI4FM'13*, HW-MACS-TR-0100. Heriot-Watt University CS, UK, 2013.

18. L. Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.
19. NASA. Formal methods, specification and verification guidebook for verification of software and computer systems. vol 2: A practitioner's companion. Technical Report NASA-GB-001-97, Washington, DC 20546, USA, May 1997.
20. I. Whiteside, D. Aspinall, L. Dixon, and G. Grov. Towards formal proof script refactoring. In *CICM*, volume 6824 of *LNCS*, pages 260–275. Springer, 2011.